

METHOD AND SYSTEM FOR PROCESSING UPSTREAM PACKETS OF AN OPTICAL NETWORK

STATEMENT REGARDING RELATED APPLICATIONS

5 This application is a continuation-in-part of a non-provisional patent application entitled, "System and Method for Communicating Optical Signals between a Data Service Provider and Subscribers," filed on July 5, 2001 and assigned U.S. Application Serial No. 09/899,410. The present application is also related to non-provisional application entitled, "System and Method for Communicating Optical
10 Signals Upstream and Downstream between a Data Service Provider and Subscribers," filed on October 4, 2001 and assigned U.S. Serial No. 09/971,363. The present application claims priority to provisional patent application entitled, "Systems to Provide Video, Voice and Data services via Fiber Optic Cable - Part 2," filed on October 26, 2000 and assigned U.S. Application Serial No. 60/244,052; provisional
15 patent application entitled, "Systems to Provide Video, Voice and Data services via Fiber Optic Cable - Part 3," filed on December 28, 2000 and assigned U.S. Application Serial No. 60/258,837; provisional patent application entitled, "Protocol to Provide Voice and Data Services via Fiber Optic Cable," filed on October 27, 2000 and assigned U.S. Application Serial No. 60/243,978; and provisional patent
20 application entitled, "Protocol to Provide Voice and Data Services via Fiber Optic Cable-Part 2," filed on May 8, 2001 and assigned U.S. Application Serial No. 60/289,112, the entire contents of each of these applications are incorporated by reference.

25 TECHNICAL FIELD

The present invention relates to video, voice, and data communication. More particularly, the present invention relates to a system and method for communicating upstream optical signals from Subscribers to a data service provider.

BACKGROUND OF INVENTION

The increasing reliance on communication networks to transmit more complex data, such as voice and video traffic, is causing a very high demand for bandwidth. To resolve this demand for bandwidth, communication networks are relying more upon optical fibers to transmit this complex data. Conventional communication architectures that employ coaxial cables are slowly being replaced with communication networks that comprise only fiber optic cables. One advantage that optical fibers have over coaxial cables is that a much greater amount of information can be carried on an optical fiber.

The Fiber-to-the-home (FTTH) optical network architecture has been a dream of many data service providers because of the aforementioned capacity of optical fibers that enable the delivery of any mix of high-speed services to businesses and consumers over highly reliable networks. Related to FTTH is fiber-to-the-business (FTTB). FTTH and FTTB architectures are desirable because of improved signal quality, lower maintenance, and longer life of the hardware involved with such systems. However, in the past, the cost of FTTH and FTTB architectures have been considered prohibitive. But now, because of the high demand for bandwidth and the current research and development of improved optical networks, FTTH and FTTB have become a reality.

A conventional hybrid fiber-to-the-home (FTTH)/hybrid fiber-coax (HFC) architecture has been proposed by the industry. HFC is currently the architecture of choice for many cable television systems. In this FTTH/HFC architecture, an active signal source is placed between the data service hub and the subscriber. Typically, in this architecture, the active source comprises a router. This conventional router typically has multiple data ports that are designed to support individual subscribers. More specifically, the conventional router uses a single port for each respective subscriber. Connected to each data port of the router is an optical fiber which, in turn, is connected to the subscriber. The connectivity between data ports and optical fibers with this conventional FTTH/HFC architecture yield a very fiber intensive last mile.

It noted that the terms, "last mile" and "first mile", are both generic terms used to

describe the last portion of an optical network that connects to subscribers. Therefore, the distance of a mile should not to be taken literally.

In addition to a high number of optical cables originating from the router, the FTTH/HFC architecture requires radio frequency signals to be propagated along traditional coaxial cables. Because of the use of coaxial cables, numerous radio frequency (RF) amplifiers are needed between the subscriber and the data service help. For example, RF amplifiers are typically needed every one to three kilometers in a coaxial type system.

The use of coaxial cables and the FTTH/HFC architecture adds to the overall cost of the system because two separate and distinct networks are present in such an architecture. In other words, the FTTH/HFC architecture has high maintenance cost because of the completely different wave guides (coaxial cable in combination with optical fiber) in addition to the electrical and optical equipment needed to support such two distinct systems. Stated more simply, the FTTH/HFC architecture merely combines an optical network with an electrical network with both networks running independently of one another.

One problem with the electrical network in the FTTH/HFC architecture involves cable modem technology which supports the data communications between the data service provider and the subscriber. The data service subscriber typically employs a cable modem termination system (CMTS) to originate downstream data communications that are destined to the subscriber. To receive these downstream data communications, the subscriber will typically use a cable modem that operates according to a particular protocol known in the industry as Data-Over-Cable-Service-Interface-Specification (DOCSIS). The DOCSIS protocol defines service flows, which are identifications assigned to groups of packets by the CMTS for the downstream flows based on an inspection of a number of parameters in a packet.

More specifically, a service flow is media access control (MAC)-layer transport service that provides unique directional transport of packets either to upstream packets transmitted by the cable modem or to downstream packets transmitted by the CMTS. The identifications assigned to groups of packets in the

DOCSIS protocol can include parameters such as TCP, UTP, IP, LLC, and 802.1 P/Q identifiers contained in an incoming packet.

- Based on these identifications, the CMTS assigns a service flow ID (SFID) to a particular datastream. A service flow typically exists when the CMTS assigns this
- 5 SFID to a datastream. The SFID serves as the principle identifier in the CMTS for the service flow. A service flow is characterized by at least an SFID and an associated direction.

- A SFID is usually assigned when a user wishes to communicate. When the user relinquishes the communications channel, the SFID ceases to exist, and if all
- 10 sessions are closed, all communications with that particular modem cease. When the modem needs to communicate again, it typically must contend for a timeslot in which is then asks for bandwidth. Thus, there will be some delay before communication can restart. This delay may or may not be noticeable to a user, depending on system loading and the nature of the application.

- 15 Communication via cable modem is asymmetrical. That is, the data rate that can be achieved in the downstream direction is greater than that which may be achieved in the upstream direction. This is adequate for certain types of communications, such as web surfing. But it is not efficient for other types of services, such as peer-to-peer file transfer (e.g., digital audio file transfers such as
- 20 Napster type services), nor is it good for large email attachments. This asymmetrical communication is a consequence of using coaxial cable with the need to restrict return signals to low frequencies.

- Accordingly, there is a need in the art for a system and method for communicating optical signals between a data service provider and a subscriber that
- 25 eliminates the use of coaxial cables and related hardware and software necessary to support the data signals propagating along the coaxial cables. There is also need in the art for a system and method for communicating optical signals between a data service provider and a subscriber that can service a large number of subscribers while reducing the number of connections at the data service hub.

There is also a need in the art for a method and system for handling upstream optical communications that can offer guaranteed bandwidth for each subscriber that is part of the optical network. There is a further need in the art for a system and method that can offer guaranteed bandwidth while using any left over portion of this guaranteed bandwidth by placing it in a pool that is accessible by each subscriber. In other words, there is a need in the art for a system and method that can reclaim guaranteed bandwidth among multiple subscribers and place this reclaimed bandwidth in a pool. Another need exists in the art for a system and method that can process group or aggregated packets with token bucket emulation. A further need in the art exists for a system and method that can provide a fair allocation of bandwidth to different subscribers where "fair" has a precise mathematical definition. Another need exists in the art for a system and method that can alleviate communication traffic load on an upstream data path.

SUMMARY OF THE INVENTION

The present invention is generally drawn to a system and method for efficient propagation of data and broadcast signals over an optical fiber network. More specifically, the present invention is generally drawn to a method and system for handling upstream optical communications originating from subscribers of an optical network that are transmitted to a data service hub of the optical network. The term "upstream" can define a communication direction where a subscriber originates a data signal that is sent upwards towards a data service hub of an optical network. Conversely, the term "downstream" can define a communication direction where a data service hub originates a data signal that is sent downwards towards subscribers of an optical network.

The method and system of the present invention can coordinate the operation of subscriber optical interfaces that are typically part of predetermined subscriber groups. The method and system can comprise a protocol that can control the time at which each optical interface is permitted to transmit data to a transceiver node. Such controlling of the time at which optical interfaces of a certain group are permitted to

transmit can be referred to as a form of Time Division Multiple Access (TDMA). With TDMA, the protocol can prevent collisions of upstream transmissions between the subscriber optical interfaces of a particular subscriber group.

- Unlike protocols and transmitters of the conventional art that have carrier sense or collision detection, the protocol and subscriber optical interfaces of the present invention are not designed to detect or sense transmissions of other optical interfaces of a subscriber group that share the same transceiver node. In other words, the transceiver node executing the protocol of the present invention controls or schedules the time periods in which a subscriber optical interface of a subscriber group is permitted to transmit information to the network.

- With the protocol of the present invention, the transceiver node can allocate additional or reduced upstream bandwidth based upon the demand of one or more subscribers. That is, the transceiver node can monitor (or police) and adjust a subscriber's upstream bandwidth on a subscription basis or on an as-needed basis.
- The transceiver node can offer upstream data bandwidth to the subscriber in preassigned increments. For example, the transceiver node can offer a particular subscriber or groups of subscribers upstream bandwidth in units of 1, 2, 5, 10, 20, 50, 100, 200, and 450 Megabits per second (Mb/s).

- One of the components of the transceiver node that executes or runs the protocol of the present invention is an optical tap routing device. The optical tap routing device can determine which optical tap multiplexer is to receive a downstream electrical signal, or identify which of the plurality of optical taps originated an upstream optical signal. The optical tap routing device can format data and implement the protocol required to send and receive data from each individual subscriber connected to a respective optical tap (as will be discussed below). The optical tap routing device can provide a form of centralized control for a group of subscribers. The optical tap routing device can comprise a computer or a hardwired apparatus that executes the program defining the inventive protocol of the present invention.

The algorithm of the present invention can account for aggregates of packets rather than individual packets. By performing calculation on aggregates of packets, the algorithm can execute less frequently which, in turn, permits its implementation in lower performance and lower cost devices. In the present invention, this approach can
5 allow implementation with software executing in a general purpose microprocessor instead of special purpose high-speed hardware.

Unlike conventional token bucket emulation that handles each data packet singly, the computer implemented protocol of the present invention can process grouped or aggregated packets with token bucket emulation. The emulation of a
10 token bucket can enable the inventive protocol to handle bandwidth very efficiently. The token bucket emulation can permit each subscriber of a subscriber group to have a guaranteed bandwidth. However, any portion of guaranteed bandwidth not used by a particular subscriber can be reclaimed by the protocol and offered to other subscribers who may need the bandwidth at a particular instant of time. The
15 reclaimed bandwidth can be offered from a pool of bandwidth that is formed for a subscriber group.

The reclaimed bandwidth, guaranteed bandwidth, and token bucket emulation can comprise functions of several parameters tracked by the inventive protocol: peak rate, sustained rate, burst size, priority weighting, and guaranteed indication (where
20 guaranteed indication tracks whether a subscriber is guaranteed a sustained rate at all times). Unlike the conventional art where the first three parameters of peak rate, sustained rate, and burst size are tracked with hardware, the present invention can monitor these parameters with the inventive protocol that can be executed in software. All of these parameters can be adjusted by the network provider to adapt the optical
25 network to various subscriber environments.

Unlike conventional routers that typically first receive a packet and then use a token bucket algorithm to decide how to handle the received packet, the present invention can be notified of the desire of the subscriber optical interface to send a group of packets and then the optical tap routing device can employ a token bucket
30 algorithm to notify the subscriber optical interface of the number of bytes that the

optical tap routing device can receive. That is, the optical tap routing device can use the token bucket algorithm to determine whether a subscriber optical interface can send (a) all of the packets it wants to send, or (b) some of the packets it wants to send, or (c) none of the packets it wants to send. If the subscriber optical interface cannot
5 send all the packets it wants to send, it can make a decision to either drop the packets or to hold them and try to send them at later time.

In this way, upstream data traffic can be managed more efficiently. Upstream information can be taken off the incoming data path: if a packet is not going to be accepted by the optical tap routing device, the subscriber optical interface doesn't
10 send it. Stated differently, the packets at the subscriber optical interface can be retained for transmission in the next transmission window (referred to later as a superframe). If the memory that retains the packet is full, no more packets are accepted from the user's equipment by the transceiver node.

Furthermore, the optical tap routing device can further make the decision of
15 whether or not to allow the subscriber optical interface to send the packet(s), based not only on the type of packet and the data load from the sending subscriber optical interface, but it also on how much data other subscriber optical interfaces in the same subscriber group want to send, and the service level to be provided to each subscriber optical interface (such as based on what amount the subscriber is paying for service,
20 or other criteria defined by the operator).

Additionally, the decision on how many packets to accept can be based on not only the total load presented by that group of subscriber optical interfaces (which share a common path coming into the optical tap routing device), but also on what is known about the total load that is going to be leaving the optical tap routing device,
25 destined for the data service hub.

In one exemplary embodiment, there can be six logical groups of subscribers, each feeding data into the optical tap routing device. There are typically one to four data paths leaving the optical tap routing device, going to the data service hub. The optical tap routing device can forecast its total load from the six logical groups, and
30 decide whether it can accept all packets and forward them on to the data service hub,

or whether it will only be able to accept a subset of the packets offered. If the optical tap routing device accepted them all, it would not be able to forward all of them to the data service hub.

While the determination of whether to accept subscriber optical interface packets can permit the inventive protocol to allocate bandwidth efficiently and with software, the present invention can also provide centralized control of upstream traffic and a fair allocation of bandwidth between groups of subscribers, where "fair" can have a mathematical definition. The mathematical definition of "fair" can comprise a weighted max-min algorithm. In this algorithm, a weight factor can comprise the product of provisioned weight and a low-pass filtered measurement of buffer occupancy for each subscriber optical interface of a group of subscriber optical interfaces. The max-min parameter for each subscriber optical interface can comprise its maximum allocation.

Max-min fairness maximizes the minimum that any subscriber is forced to accept. In other words, the result of a max-min allocation is that the subscriber that receives the smallest allocation is assured that he could not have been given a greater allocation without some other subscriber receiving even less than he did. A weighed max-min allocation assigns each party a relative weighting factor. It maximizes the minimum of that weighting factor multiplied by the allocation, and provides a rigorous way to bias the max-min allocation in favor of parties with more priority.

BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 is a functional block diagram of some of the core components of an exemplary optical network architecture according to the present invention.

Fig. 2 is a functional block diagram illustrating an exemplary optical network architecture for the present invention.

Fig. 3 is a functional block diagram illustrating an exemplary transceiver node according to the present invention.

Fig. 4 is a block diagram illustrating an optical tap connected to a subscriber interface by a single optical waveguide according to one exemplary embodiment of the present invention.

Fig. 5 is a block diagram illustrating guaranteed upstream bandwidth and pooled upstream bandwidth during a superframe according to a first exemplary embodiment of the present invention.

Fig. 6 is a block diagram illustrating another superframe comprising variable length time slots for subscriber optical interfaces of a predetermined group of sixteen subscribers according to a second exemplary embodiment of the present invention.

Fig. 7 is a block diagram illustrating a group of data packets transmitted during a single time slot according to the second exemplary embodiment of the present invention.

Fig. 8 is a block diagram illustrating exemplary contents and field sizes of a data packet according to the second exemplary embodiment of the present invention.

Fig. 9 is a block diagram illustrating multiple superframes according to the second exemplary embodiment of the present invention.

Fig. 10 is a block diagram illustrating multiple superframes and possible queuing and transmission delays according to the second exemplary embodiment of the present invention.

Fig. 11 is a logic flow diagram illustrating one exemplary embodiment of the high level steps that the invention performs in managing upstream data.

Fig. 13 is a logic flow diagram illustrating a more detailed exemplary subprocess of step 1220 in Fig. 11.

Fig. 14 is a logic flow diagram illustrating a more detailed exemplary subprocess of step 1325 in Fig. 12.

Fig. 15 is a logic flow diagram illustrating a more detailed exemplary subprocess of step 1330 in Fig. 12.

Fig. 16 is a logic flow diagram illustrating a more detailed exemplary subprocess of step 1335 in Fig. 12.

Fig. 17 is a logic flow diagram illustrating a more detailed exemplary subprocess of step 1365 in Fig. 12.

5

DETAILED DESCRIPTION OF EXEMPLARY EMBODIMENTS

The present invention may be embodied in software or hardware or a combination thereof disposed within an optical network. The present invention can comprise a transceiver node that further comprises an optical tap routing device and a plurality of optical tap multiplexers for receiving upstream packets from one or more subscriber optical interfaces. Unlike conventional token bucket emulation that handles data packets in a single manner, the protocol of the present invention can process aggregates of data packets with token bucket emulation. With the protocol of the present invention, a transceiver node can allocate additional or reduced upstream bandwidth based upon the demand of one or more subscribers. The present invention can monitor and adjust a subscriber's upstream bandwidth on a subscription basis or on an as-needed basis. The present invention can permit each subscriber of a subscriber group to have a guaranteed bandwidth. However, most portions of guaranteed bandwidth not used by a particular subscriber can be reclaimed by the protocol and offered to other subscribers who may need the bandwidth at a particular instant of time.

Referring now to the drawings in which like numerals represent like elements throughout the several figures, aspects of the present invention and the illustrative operating environment will be described.

Figure 1 is a functional block diagram illustrating an exemplary optical network architecture 100 according to the present invention. The exemplary optical network architecture 100 comprises a data service hub 110 that is connected to one or more transceiver nodes 120. The transceiver nodes 120, in turn, are connected to an optical taps 130. The term "transceiver node" can refer to nodes as described in compending application titled, "System and Method for Communicating Optical Signals

Between A Data Service Provider and Subscribers,” Serial No. 09/899,410, filed on July 5, 2001 and copending application titled, “System and Method for Communicating Optical Signals Upstream and Downstream between a Data Service Provider and Subscribers,” filed on October 4, 2001 and assigned U.S. Serial No. 09/971,363, the contents of these two applications are herein incorporated by reference.

The optical taps 130 can be connected to a plurality of subscriber optical interfaces 140. Between respective components of the exemplary optical network architecture 100 are optical waveguides such as optical waveguides 150, 160, 170, and 180. The optical waveguides 150-180 are illustrated by arrows where the arrowheads of the arrows illustrate exemplary directions of data flow between respective components of the illustrative and exemplary optical network architecture 100. While only an individual transceiver node 120, an individual optical tap 130, and an individual subscriber optical interface 140 are illustrated in Figure 1, as will become apparent from Figure 2 and its corresponding description, a plurality of transceiver nodes 120, optical taps 130, and subscriber optical interfaces 140 can be employed without departing from the scope and spirit of the present invention. Typically, in many of the exemplary embodiments of the present invention, multiple subscriber optical interfaces 140 are connected to one or more optical taps 130.

The transceiver node 120 can allocate additional or reduced bandwidth based upon the demand of one or more subscribers that use the subscriber optical interfaces 140. The transceiver node 120 can be designed to withstand outdoor environmental conditions and can be designed to hang on a strand or fit in a pedestal or “hand hole.” The transceiver node can operate in a temperature range between minus 40 degrees Celsius to plus 60 degrees Celsius. The transceiver node 120 can operate in this temperature range by using passive cooling devices that do not consume power.

Unlike the conventional routers disposed between the subscriber optical interface 140 and data service hub 110, the transceiver node 120 does not require active cooling and heating devices that control the temperature surrounding the transceiver node 120. The present invention places more of the decision-making

electronics at the transceiver node 120 rather than at the subscriber optical interface 140. Typically, the decision-making electronics are larger in size and more costly than the electronics placed in the subscriber optical interface 140 of the present invention.

- 5 Because the transceiver node 120 does not require active temperature controlling devices, the transceiver node 120 lends itself to a compact electronic packaging volume that is typically smaller than the environmental enclosures of conventional routers.

- 10 In one exemplary embodiment of the present invention, three trunk optical waveguides 160, 170, and 180 (that can comprise optical fibers) can conduct optical signals from the data service hub 110 to the transceiver node 120. It is noted that the term "optical waveguide" used in the present application can apply to optical fibers, planar light guide circuits, and fiber optic pigtails and other like optical waveguides.

- 15 A first optical waveguide 160 can carry broadcast video and other signals. The signals can be carried in a traditional cable television format wherein the broadcast signals are modulated onto carriers, which in turn modulate an optical transmitter (not shown) in the data service hub 110. A second optical waveguide 170 can carry downstream targeted services such as data and telephone services to be delivered to one or more subscriber optical interfaces 140. In addition to carrying
20 subscriber-specific optical signals, the second optical waveguide 170 can also propagate internet protocol broadcast packets, as is understood by those skilled in the art.

- In one exemplary embodiment, a third optical waveguide 180 can transport data signals upstream from the transceiver node 120 to the data service hub 110. The
25 optical signals propagated along the third optical waveguide 180 can also comprise data and telephone services received from one or more subscribers. Similar to the second optical waveguide 170, the third optical waveguide 180 can also carry IP broadcast packets, as is understood by those skilled in the art.

- 30 The third or upstream optical waveguide 180 is illustrated with dashed lines to indicate that it is merely an option or part of one exemplary embodiment according to

the present invention. In other words, the third optical waveguide 180 can be removed. In another exemplary embodiment, the second optical waveguide 170 propagates optical signals in both the upstream and downstream directions as is illustrated by the double arrows depicting the second optical waveguide 170. In such an exemplary embodiment where the second optical waveguide 170 propagates bidirectional optical signals, only two optical waveguides 160, 170 would be needed to support the optical signals propagating between the data server's hub 110 in the transceiver node 120. In another exemplary embodiment (not shown), a single optical waveguide can be the only link between the data service hub 110 and the transceiver node 120. In such a single optical waveguide embodiment, three different wavelengths can be used for the upstream and downstream signals. Alternatively, bi-directional data could be modulated on one wavelength.

In one exemplary embodiment, the optical tap 130 can comprise an 8-way optical splitter. This means that the optical tap 130 comprising an 8-way optical splitter can divide downstream optical signals eight ways to serve eight different subscriber optical interfaces 140. In the upstream direction, the optical tap 130 can combine the optical signals received from the eight subscriber optical interfaces 140.

In another exemplary embodiment, the optical tap 130 can comprise a 4-way splitter to service four subscriber optical interfaces 140. Yet in another exemplary embodiment, the optical tap 130 can further comprise a 4-way splitter that is also a pass-through tap meaning that a portion of the optical signal received at the optical tap 130 can be extracted to serve the 4-way splitter contained therein while the remaining optical energy is propagated further downstream to another optical tap or another subscriber optical interface 140. The present invention is not limited to 4-way and 8-way optical splitters. Other optical taps having fewer or more than 4-way or 8-way splits are not beyond the scope of the present invention.

Referring now to Figure 2, this Figure is a functional block diagram illustrating an exemplary optical network architecture 100 that further includes subscriber groupings 200 that correspond with a respective transceiver node 120. Figure 2 illustrates the diversity of the exemplary optical network architecture 100

where a number of optical waveguides 150 connected between the transceiver node 120 and the optical taps 130 is minimized. Figure 2 also illustrates the diversity of subscriber groupings 200 that can be achieved with the optical tap 130.

Each optical tap 130 can comprise an optical splitter. The optical tap 130 allows multiple subscriber optical interfaces 140 to be coupled to a single optical waveguide 150 that is connected to the transceiver node 120. In one exemplary embodiment, six optical fibers 150 are designed to be connected to the transceiver node 120. Through the use of the optical taps 130, sixteen subscribers can be assigned to each of the six optical fibers 150 that are connected to the transceiver node 120.

In another exemplary embodiment, twelve optical fibers 150 can be connected to the transceiver node 120 while eight subscriber optical interfaces 140 are assigned to each of the twelve optical fibers 150. Those skilled in the art will appreciate that the number of subscriber optical interfaces 140 assigned to a particular waveguide 150 that is connected between the transceiver node 120 and a subscriber optical interface 140 (by way of the optical tap 130) can be varied or changed without departing from the scope and spirit of the present invention. Further, those skilled in the art recognize that the actual number of subscriber optical interfaces 140 assigned to the particular fiber optic cable is dependent upon the amount of power available on a particular optical fiber 150.

As depicted in subscriber grouping 200, many configurations for supplying communication services to subscribers are possible. For example, while optical tap 130_A can connect subscriber optical interfaces 140_{A1} through subscriber optical interface 140_{AN} to the laser transmitter node 120, optical tap 130_A can also connect other optical taps 130 such as optical tap 130_{AN} to the transceiver node 120. The combinations of optical taps 130 with other optical taps 130 in addition to combinations of optical taps 130 with subscriber optical interfaces 140 are limitless. With the optical taps 130, concentrations of distribution optical waveguides 150 at the transceiver node 120 can be reduced. Additionally, the total amount of fiber needed to service a subscriber grouping 200 can also be reduced.

With the active transceiver node 120 of the present invention, the distance between the transceiver node 120 and the data service hub 110 can comprise a range between 0 and 80 kilometers. However, the present invention is not limited to this range. Those skilled in the art will appreciate that this range can be expanded by selecting various off-the-shelf components that make up several of the devices of the present system.

Those skilled in the art will appreciate that other configurations of the optical waveguides disposed between the data service hub 110 and transceiver node 120 are not beyond the scope of the present invention. Because of the bi-directional capability of optical waveguides, variations in the number and directional flow of the optical waveguides disposed between the data service hub 110 and the transceiver node 120 can be made without departing from the scope and spirit of the present invention.

Those skilled in the art will appreciate that the selection of optical transmitters comprising optical waveguide transceiver 430 and data service hub 110 for each circuit may be optimized for the optical path lengths needed between the data service hub 110 and the transceiver node 120. Further, those skilled in the art will appreciate that the wavelengths discussed are practical but are only illustrative in nature. In some scenarios, it may be possible to use communication windows at 1310 and 1550 nm in different ways without departing from the scope and spirit of the present invention. Further, the present invention is not limited to a 1310 and 1550 nm wavelength regions. Those skilled in the art will appreciate that smaller or larger wavelengths for the optical signals are not beyond the scope and spirit of the present invention.

Referring now to Figure 3, this Figure illustrates a functional block diagram of an exemplary transceiver node 120 of the present invention. In this exemplary embodiment, the transceiver node 120 can comprise a unidirectional optical signal input port 405 that can receive optical signals propagated from the data service hub 110 that are propagated along a first optical waveguide 160. The optical signals received at the unidirectional optical signal input port 405 can comprise broadcast video data. The optical signals received at the input port 405 are propagated to an

amplifier 410 such as an Erbium Doped Fiber Amplifier (EDFA) in which the optical signals are amplified. The amplified optical signals are then propagated to a splitter 415 that divides the broadcast video optical signals among diplexers 420 that are designed to forward optical signals to predetermined subscriber groups 200.

5 The transceiver node 120 can further comprise a bi-directional optical signal input/output port 425 that connects the transceiver node 120 to a second optical waveguide 170 that supports bi-directional data flow between the data service hub 110 and transceiver node 120. Downstream optical signals flow through the bi-directional optical signal input/output port 425 to an optical waveguide transceiver 430 that
10 converts downstream optical signals into the electrical domain. The optical waveguide transceiver further converts upstream electrical signals into the optical domain. The optical waveguide transceiver 430 can comprise an optical/electrical converter and an electrical/optical converter. In another exemplary system, fiber 170 propagates signals downstream and fiber 180 propagates signals upstream.

15 Downstream and upstream electrical signals are communicated between the optical waveguide transceiver 430 and an optical tap routing device 435. The optical tap routing device 435 can manage the interface with the data service hub optical signals and can route or divide or apportion the data service hub signals according to individual tap multiplexers 440 that communicate optical signals with one or more
20 optical taps 130 and ultimately one or more subscriber optical interfaces 140. It is noted that tap multiplexers 440 operate in the electrical domain to modulate laser transmitters in order to generate optical signals that are assigned to groups of subscribers coupled to one or more optical taps.

 The optical tap routing device 435 is notified of available upstream data
25 packets as they arrive, by each tap multiplexer 440. The optical tap routing device is connected to each tap multiplexer 440 to receive these upstream data packets. The optical tap routing device 435 relays the packets to the data service hub 110 via the optical waveguide transceiver 430. The optical tap routing device 435 can build a lookup table from these upstream data packets coming to it from all tap multiplexers
30 440 (or ports), by reading the source IP address of each packet, and associating it with

the tap multiplexer 440 through which it came. This lookup table can then be used to route packets in the downstream path. As each packet comes in from the optical waveguide transceiver 430, the optical tap routing device 435 looks at the destination IP address (which is the same as the source IP address for the upstream packets).

- 5 From the lookup table the optical tap routing device 435 can determine which port is connected to that IP address, so it sends the packet to that port. This can be described as a normal layer 3 router function as is understood by those skilled in the art.

- The optical tap routing device 435 can assign multiple subscribers to a single port. More specifically, the optical tap routing device 435 can service groups of
10 subscribers with corresponding respective, single ports. The optical taps 130 coupled to respective tap multiplexers 440 can supply downstream optical signals to pre-assigned groups of subscribers who receive the downstream optical signals with the subscriber optical interfaces 140.

- In other words, the optical tap routing device 435 can determine which tap
15 multiplexer 440 is to receive a downstream electrical signal, or identify which of a plurality of optical taps 130 propagated an upstream optical signal (that is converted to an electrical signal). The optical tap routing device 435 can format data and implement the protocol required to send and receive data from each individual subscriber connected to a respective optical tap 130. The optical tap routing device
20 435 can comprise a computer or a hardwired apparatus that defines a protocol for communications with groups of subscribers assigned to individual ports.

- The single ports of the optical tap routing device are connected to respective tap multiplexers 440. With the optical tap routing device 435, the transceiver node 120 can adjust a subscriber's bandwidth on a subscription basis or on
25 an as-needed or demand basis. The transceiver node 120 via the optical tap routing device 435 can offer data bandwidth to subscribers in pre-assigned increments. For example, the transceiver node 120 via the optical tap routing device 435 can offer a particular subscriber or groups of subscribers upstream and downstream bandwidth in units of 1, 2, 5, 10, 20, 50, 100, 200, and 450 Megabits per second (Mb/s). Those

skilled in the art will appreciate that other subscriber bandwidth units are not beyond the scope of the present invention.

Electrical signals are communicated between the optical tap routing device 435 and respective tap multiplexers 440. The tap multiplexers 440 propagate optical signals to and from various groupings of subscribers. Each tap multiplexer 440 is connected to a respective optical transmitter 325. Each optical transmitter 325 can comprise one of a Fabry-Perot (F-P) laser, a distributed feedback laser (DFB), or a Vertical Cavity Surface Emitting Laser (VCSEL). However, other types of optical transmitters are possible and are not beyond the scope of the present invention. The optical transmitters produce the downstream optical signals that are propagated towards the subscriber optical interfaces 140.

Each tap multiplexer 440 is also coupled to an optical receiver 370. From the bi-directional splitter 360, respective optical receivers 370 can convert the upstream optical signals into the electrical domain. Each optical receiver 370 can comprise one or more photoreceptors or photodiodes that convert optical signals into electrical signals. Since the optical transmitters 325 and optical receivers 370 can comprise off-the-shelf hardware to generate and receive respective optical signals, the transceiver node 120 lends itself to efficient upgrading and maintenance to provide significantly increased data rates.

Each optical transmitter 325 and each optical receiver 370 are connected to a respective bi-directional splitter 360. Each bi-directional splitter 360 in turn is connected to a diplexer 420 which combines the unidirectional optical signals received from the splitter 415 with the downstream optical signals received from respective optical receivers 370. In this way, broadcast video services as well as data services can be supplied with a single optical waveguide such as a distribution optical waveguide 150 as illustrated in Figure 2. In other words, optical signals can be coupled from each respective diplexer 420 to a combined signal input/output port 445 that is connected to a respective distribution optical waveguide 150.

Unlike the conventional art, the transceiver node 120 does not employ a conventional router. The components of the transceiver node 120 can be disposed

within a compact electronic packaging volume. For example, the transceiver node 120 can be designed to hang on a strand or fit in a pedestal similar to conventional cable TV equipment that is placed within the "last mile," or subscriber-proximate portions of a network. It is noted that the term, "last mile," is a generic term often used to describe the last portion of an optical network that connects to subscribers.

Also because the optical tap routing device 435 is not a conventional router, it does not require active temperature controlling devices to maintain the operating environment at a specific temperature. In other words, the transceiver node 120 can operate in a temperature range between minus 40 degrees Celsius to 60 degrees Celsius in one exemplary embodiment.

While the transceiver node 120 does not comprise active temperature controlling devices that consume power to maintain temperature of the transceiver node 120 at a single temperature, the transceiver node 120 can comprise one or more passive temperature controlling devices 450 that do not consume power. The passive temperature controlling devices 450 can comprise one or more heat sinks or heat pipes that remove heat from the transceiver node 120. Those skilled in the art will appreciate that the present invention is not limited to these exemplary passive temperature controlling devices. Further, those skilled in the art will also appreciate the present invention is not limited to the exemplary operating temperature range disclosed. With appropriate passive temperature controlling devices 450, the operating temperature range of the transceiver node 120 can be reduced or expanded.

In addition to the transceiver node's 120 ability to withstand harsh outdoor environmental conditions, the transceiver node 120 can also provide high speed symmetrical data transmissions. In other words, the transceiver node 120 can propagate the same bit rates downstream and upstream to and from a network subscriber. This is yet another advantage over conventional networks, which typically cannot support symmetrical data transmissions as discussed in the background section above. Further, the transceiver node 120 can also serve a large number of subscribers while reducing the number of connections at both the data service hub 110 and the transceiver node 120 itself.

The transceiver node 120 also lends itself to efficient upgrading that can be performed entirely on the network side or data service hub 110 side. That is, upgrades to the hardware forming the transceiver node 120 can take place in locations between and within the data service hub 110 and the transceiver node 120. This means that the subscriber side of the network (from distribution optical waveguides 150 to the subscriber optical interfaces 140) can be left entirely in-tact during an upgrade to the transceiver node 120 or data service hub 110 or both.

Referring now to Figure 4, this Figure is a functional block diagram illustrating an optical tap 130 connected to a subscriber optical interface 140 by a single optical waveguide 150 according to one exemplary embodiment of the present invention. The optical tap 130 can comprise a combined signal input/output port 505 that is connected to another distribution optical waveguide that is connected to a transceiver node 120. As noted above, the optical tap 130 can comprise an optical splitter 510 that can be a 4-way or 8-way optical splitter. Other optical taps having fewer or more than 4-way or 8-way splits are not beyond the scope of the present invention. The optical tap can divide downstream optical signals to serve respective subscriber optical interfaces 140. In the exemplary embodiment in which the optical tap 130 comprises a 4-way optical tap, such an optical tap can be of the pass-through type, meaning that a portion of the downstream optical signals is extracted or divided to serve a 4-way splitter contained therein, while the rest of the optical energy is passed further downstream to other distribution optical waveguides 150.

The optical tap 130 is an efficient coupler that can communicate optical signals between the transceiver node 120 and a respective subscriber optical interface 140. Optical taps 130 can be cascaded, or they can be connected in a star architecture from the transceiver node 120. As discussed above, the optical tap 130 can also route signals to other optical taps that are downstream relative to a respective optical tap 130.

The optical tap 130 can also connect to a limited or small number of optical waveguides so that high concentrations of optical waveguides are not present at any particular transceiver node 120. In other words, in one exemplary embodiment, the

optical tap can connect to a limited number of optical waveguides 150 at a point remote from the transceiver node 120 so that high concentrations of optical waveguides 150 at a transceiver node can be avoided. However, those skilled in the art will appreciate that the optical tap 130 can be incorporated within the transceiver
5 node 120 as will be discussed in further detail below with respect to another exemplary embodiment of the transceiver node 120 as illustrated in Figure 11.

The subscriber optical interface 140 functions to convert downstream optical signals received from the optical tap 130 into the electrical domain that can be processed with appropriate communication devices. The subscriber optical interface
10 140 further functions to convert upstream electrical signals into upstream optical signals that can be propagated along a distribution optical waveguide 150 to the optical tap 130. The subscriber optical interface 140 can comprise an optical diplexer 515 that divides the downstream optical signals received from the distribution optical waveguide 150 between a bi-directional optical signal splitter 520 and an analog
15 optical receiver 525. The optical diplexer 515 can receive upstream optical signals generated by a digital optical transmitter 530. The digital optical transmitter 530 converts electrical binary/digital signals to optical form so that the optical signals can be transmitted back to the data service hub 110. Conversely, the digital optical receiver 540 converts optical signals into electrical binary/digital signals so that the
20 electrical signals can be handled by processor 550.

The present invention can propagate the optical signals at various wavelengths. However, the wavelength regions discussed are practical and are only illustrative of exemplary embodiments. Those skilled in the art will appreciate that other wavelengths that are either higher or lower than or between the 1310 and 1550
25 nm wavelength regions are not beyond the scope of the present invention.

The analog optical receiver 525 can convert the downstream broadcast optical video signals into modulated RF television signals that are propagated out of the modulated RF unidirectional signal output 535. The modulated RF unidirectional signal output 535 can feed to RF receivers such as television sets (not shown) or
30 radios (not shown). The analog optical receiver 525 can process analog modulated RF

transmission as well as digitally modulated RF transmissions for digital TV applications.

The bi-directional optical signal splitter 520 can propagate combined optical signals in their respective directions. That is, downstream optical signals entering the bi-directional optical splitter 520 from the optical the optical diplexer 515, are propagated to the digital optical receiver 540. Upstream optical signals entering it from the digital optical transmitter 530 are sent to optical diplexer 515 and then to optical tap 130. The bi-directional optical signal splitter 520 is connected to a digital optical receiver 540 that converts downstream data optical signals into the electrical domain. Meanwhile the bi-directional optical signal splitter 520 is also connected to a digital optical transmitter 530 that converts upstream electrical signals into the optical domain.

The digital optical receiver 540 can comprise one or more photoreceptors or photodiodes that convert optical signals into the electrical domain. The digital optical transmitter can comprise one or more lasers such as the Fabry-Perot (F-P) Lasers, distributed feedback lasers, and Vertical Cavity Surface Emitting Lasers (VCSELs).

The digital optical receiver 540 and digital optical transmitter 530 are connected to a processor 550 that selects data intended for the instant subscriber optical interface 140 based upon an embedded address. The data handled by the processor 550 can comprise one or more of telephony and data services such as an Internet service. The processor 550 is connected to a telephone input/output 555 that can comprise an analog interface. The processor 550 is also connected to a data interface 560 that can provide a link to computer devices, set top boxes, ISDN phones, and other like devices. Alternatively, the data interface 560 can also provide a link to a Voice over Internet Protocol (VoIP) telephone or Ethernet telephone. The data interface 560 can comprise one of Ethernet's (10BaseT, 100BaseT, Gigabit) interface, HPNA interface, a universal serial bus (USB) an IEEE1394 interface, an ADSL interface, and other like interfaces.

Referring now to Figure 5, this figure is a block diagram illustrating guaranteed upstream bandwidth and pooled upstream bandwidth during a single

superframe 562 according to an exemplary embodiment of the present invention. A superframe 562 typically comprises a duration of time over which an entire sequence of events repeats. In one exemplary embodiment, a superframe may comprise 8 milliseconds in duration. This means, that every 8 milliseconds, a transceiver node

5 120 can cycle through all the subscribers 140 on a given tap multiplexer 440, giving each subscriber 140 a turn or an opportunity to transmit data packets. In one exemplary embodiment, sixteen subscribers 140 can be serviced on a single upstream channel. However, those skilled in the art will appreciate that fewer or greater number of subscribers 140 can be added to channels without departing from the scope

10 and spirit of the present invention.

Adjacent to frame 500 is a horizontal axis 576 that denotes time. In Figure 5, each time slot represents a period of time in which a subscriber optical interface (SOI) can transmit data packets in an upstream direction. Time slots 564, 566, 568, 570 and 572 form the guaranteed bandwidth 574 for superframe 562. Bandwidth, as described

15 in Figure 5, comprises the number of bits of data that can be transmitted per unit of time. This is typically a function of the data rate, and also the time allowed. Bandwidth is often referred to in terms of a certain number of megabits per second (million bits per second, MBps).

As illustrated in Figure 5, each subscriber optical interface 140 has an assigned

20 time slot 564, 566, 568, 570, and 572 in which to transmit information in the upstream direction over the optical network 100. For example, a first subscriber optical interface assigned to a first time slot 564 may first transmit while a second subscriber optical interface may transmit during a second time slot 566 after the duration of time slot 564. After each subscriber optical interface has transmitted during its assigned

25 time slot, each subscriber optical interface may further transmit during a pool bandwidth 578. Each subscriber optical interface 140 typically must request bandwidth from the pool bandwidth 578 before bandwidth can be granted it by the protocol of the present invention that is executed primarily in the transceiver node 120.

Any subscriber optical interface 140 may use the pool bandwidth 578 if the pool bandwidth 578 has not been previously taken by another subscriber optical interface. Typically, a subscriber optical interface 140 can use guaranteed bandwidth first, and then if additional data is desired to be transmitted beyond that of the guaranteed bandwidth, the subscriber optical interface 140 can request for additional time that may be allocated from the pool bandwidth 578.

Referring now to Figure 6, this figure illustrates a second exemplary and preferred embodiment of how upstream bandwidth is divided into individual time slots within a single superframe. A superframe 600 typically comprises a duration of time over which an entire sequence of events repeats. In one exemplary embodiment, a superframe may comprise eight (8) milliseconds in duration. This means that every eight (8) milliseconds a transceiver node can cycle through all the subscribers on a given channel, giving each subscriber a turn or opportunity to transmit data packets. In one exemplary embodiment, sixteen subscribers can be serviced on a single upstream channel. However, those skilled in the art will appreciate that fewer or greater number of subscribers can be added to channels without departing from the scope and spirit of the present invention.

Figure 6 represents time as increasing from left to right. Each time slot represents a period of time in which a specific subscriber optical interface 140 (SOI) can transmit data packets in an upstream direction. For example, a first subscriber optical interface assigned to time slot 0 may transmit during that time slot, while a second subscriber optical interface may transmit during a second time slot after the duration of the first time slot.

The duration of each individual time slot is controlled by the transceiver node 120. The transceiver node controls each time slot so as to provide the guaranteed bandwidth to each subscriber, as well as any pooled bandwidth made available to that subscriber for the duration of the current superframe. The sum of all time slots usually must be less than or equal to the duration of the superframe. If the entire superframe is not needed to transmit the available data, then the time slots are "bunched up" near

the beginning of the superframe, and there will be time near the end of the superframe in which nothing (no data) is transmitted.

The final time slot in Figure 6 is a special time slot reserved for unknown subscriber optical interfaces to transmit, allowing the transceiver node to discover their presence. Because the transceiver node and any unknown subscriber optical interface will not have established ranging information and other parameters that may influence the relative timing of the two systems, this special discovery time slot must be of sufficient duration to accommodate a wide range of relative timing offsets. An exemplary embodiment of the present invention may establish a minimum discovery slot time of 180 microseconds.

The discovery time slot need not be present in every superframe. A typical embodiment of the present invention would have the transceiver node schedule a discovery time slot approximately once every second, or once every 125 superframes.

Figure 6 further shows that a guard band or time may be provided between adjacent time slots. The guard band provides sufficient time for the laser of the first subscriber optical interface 140 to reduce its output power to zero and the laser of the second subscriber optical interface 140 to reach full operating power without interfering with the first one. It also includes time for the tap multiplexer 440 to lock to (or synchronize with) the clock of the starting subscriber optical interface 140. A guard time of 2 μ s is shown by way of example. The exact duration of the guard band will depend on the transition times of the digital optical transmitters 530 and the clock acquisition time of the tap multiplexer 440.

The guard time typically exists, but it may be made smaller in the case of two laser optical receivers being shared at the transceiver node. In this case, if the SOI 140 ending its time slot and the SOI 140 starting its time slot, feed different optical receivers (i.e., when a switch must alternate positions to receive the starting SOI 140), then the starting SOI's optical transmitter can be turned on before the ending SOI's transmitter is turned off, and the ending SOI's transmitter can be turned off after the switch is activated. Such an operation could save some guard time. However, some guard time may exist during which the starting SOI 140 is transmitting a "run-in

sequence” required to lock (or synchronize with) the clock in the tap multiplexer to the clock in the transmitter. This is a common requirement as understood by those skilled in the art.

- Figure 7 illustrates how each individual subscriber optical interface 140 may make use of its assigned time slot. The figure shows that the subscriber optical interface 140 first transmits an SOI status message and then follows that message with zero or more data messages. The number of data messages is limited only by the duration of the assigned time slot; transmission of the entire data message must conclude before the time slot ends. Each message (both status and data) is formatted as a standard Ethernet frame (as illustrated in Figure 8), and the standard Ethernet interframe gap (such as 96 nanoseconds at 1 Gbit/s) separates individual frames. The interframe gap is understood by those skilled in the art as helping the receiver identify the start of an Ethernet message.

Table 1. - Packet Fields

Field Description	Length in Bits	Length in Bytes	Offset in Bytes
Destination address (fixed value, not used)	48	6	0
Source address (fixed value, not used)	48	6	6
Type Field, must match []Transceiver Node Type Field	16	2	12
Reserved (not used)	16	2	14
Reserved (not used)	16	2	16
Reserved (not used)	16	2	18
TN Transmission time stamp (T2), ms, 4 bits used	16	2	20
TN Transmission time stamp (T2), ls, units = 8 ns	16	2	22
Status word			
Reserved	19		
Subscriber optical interface status	5		
Subscriber optical interface id	8	4	24
Lo priority buffered data count, units = 128 bytes	16	2	28
Hi priority buffered data count, units = 128 bytes	16	2	30
Software info 0	16	2	32
Software info 1	16	2	34
Software info 2	16	2	36
Software info 3	16	2	38
Reserved (not used)	8*20	20	40
FCS	32	4	60

The SOI status message is also formatted as a standard Ethernet frame. Its payload can comprise a count of the number of bytes stored in the subscriber optical interface's upstream transmission queues at the beginning of the timeslot. This information will be used to compute how much time that SOI gets to transmit in a later superframe, as shown below. Table 1 shows an exemplary SOI status message.

Note that one time slot, from one subscriber optical interface (SOI) 140, may take up the majority of the superframe, or it may take up only a minor portion of the superframe. If only one SOI 140 has significant data to transmit, then it will be assigned nearly the entire superframe. In the preferred embodiment, each SOI has a minimum time during which it transmits its SOI status message 710. By way of example, the minimum time may be that required to transmit 256 kb of data is 256 μ s. An exception would be time slot 16,650, which may be shorter because only a small message must be transmitted. These are examples only, and the invention may use longer or shorter minimum time slots.

Referring now to Figure 8, the contents of an exemplary timeslot 800 are illustrated. Those skilled in the art will recognize this as a standard Ethernet packet, though the invention is not limited to Ethernet packets. Exemplary timeslot 800 may comprise a preamble 810 that may have an exemplary size of eight bytes. The timeslot 800 may further comprise a destination media access control message 815 that may comprise an exemplary length of six bytes. The timeslot 800 may further comprise a source media access control message 820 that also comprises an exemplary size of six bytes. Adjacent to the source media access control message 820 can be a type/length identifier, T 825, which is understood by those skilled in the art.

Adjacent to the T 825 can be message data 830 comprising an exemplary length between forty-six and 1,500 bytes. Adjacent to the message data 830 is an FCS 835. The FCS is a frame check sequence, a standard feature of Ethernet used to validate the integrity of the data in the packet. This is understood by those skilled in the art. The FCS message 840 can comprise an exemplary length of four bytes. Those skilled in the art understand that the present invention is not limited to the

exemplary lengths illustrated. That is, smaller or larger data messages are not beyond the scope of the present invention.

Referring to Figure 9, this Figure illustrates a plurality of exemplary superframes 910, 915 and 920 that can be processed by a transceiver node 120 according to the present invention. Figure 9 shows how the subscriber optical interfaces 140 notify the transceiver node 120 of their respective needs for bandwidth, and how the transceiver node 120 authorizes the subscriber optical interfaces 140 to transmit that data. Within the transceiver node 120, optical tap routing device 435, working with the plurality of tap multiplexers 440, performs the intelligent operations required to manage traffic. During superframe_{j-2} 910 the subscriber optical interfaces can make their resource reservation requests.

After the completion of superframe_{j-2} 910 and a Reservation Collection Delay (RCD) 925, the complete set of reservation requests is available to the optical tap routing device 435. The RCD 925 includes the one-way propagation delay and other and other delays within the transceiver node 120. The optical tap routing device 435 usually must complete its calculation of timeslot allocations for superframe_j 920 before the beginning of that superframe, with enough time remaining to communicate the allocations to all of the subscriber optical interfaces 140.

The Reservation Transmission Delay (RTD) 935 typically comprises the required time remaining to communicate the allocations to all of the subscriber optical interfaces 140. It can include the setup time in the optical tap routing device 435, worst-case transmission delay for the downstream message in progress (which message must complete before the setup data can be transmitted), one-way propagation delay, and the reception and interpretation delay in the subscriber optical interface 140.

Between the RCD 925 and the RTD 935 is the actual processing time 930 available for calculating slot allocations for superframe_j 920. During this time, the optical tap routing device 435 usually must make all the computations required to apportion data transmission during superframe_j 930. Note that the RCD 925, the calculation time 930 and the RTD 935 typically take place during superframe_{j-1} 915.

During superframe_{j,1} 915 data is also being transmitted upstream, based on requests made in the superframe that preceded superframe_{j,2} 910. Had this superframe been shown, it would have appeared in Figure 9 to the left of superframe_{j,2} 910.

Thus, the process illustrated in Figure 9 may be thought of as a continuous, or sliding, process, whereby during each frame data is transmitted, along with requests for transmission slots that will be granted (or not, depending on time available) to superframes in the future. At least one superframe 915 usually must intervene between the time of the requests (during superframe 910) and the time of the grant (acted on during superframe 920), to allow the optical tap routing device 435 time to gather all requests and determine how to dispense with them, and then to send appropriate instructions to each subscriber optical interface 140. However, during the intervening superframe 915, upstream data is being transmitted, based on an earlier cycle of request, decision and grant.

Referring now to Figure 10, this Figure illustrates a plurality of superframes 1010, 1015, 1020 and 1025 that have an exemplary length of τ (tau). This figure is included to illustrate the maximum delay in transmitting a packet, that results from the method used to determine what messages can be sent. It illustrates an exemplars maximum delay that a packet will encounter from the time it enters the subscriber optical interface 140, to the time it arrives at the optical tap routing device 435. From the optical tap routing device 435, the packet is handled according to the standard gigabit Ethernet protocol, defined by 802.3z, in the exemplary system. However, the invention is not limited to handling the packet using Ethernet protocol from the optical tap routing device 435 upstream into standard routers. Other known protocols such as SONET may be used at this stage of processing.

As a consequence of the procedure described above with reference to Figure 9, there is usually a minimum delay between the time when a subscriber optical interface 140 recognizes the need for resources and the time those resources are available. As Figure 10 illustrates, the worst-case value for this delay is approximately four times the length of the superframe, such as eight (8) ms in the exemplary system. Thus, in

the exemplary system, a packet may be delayed by nearly 32 ms. This worst-case delay occurs when the situation illustrated in Figure 10 occurs.

In this example, traffic 1035 arrives at the subscriber optical interface 140 just after that subscriber optical interface 140 has transmitted its data for superframe₃ 1010. The subscriber optical interface 140 cannot request a timeslot for the traffic 1035 to be transmitted, until its timeslot 1040 in superframe₂ 1015.

As shown above in Figure 9, the packet is then delayed until superframe₁ 1025, when it is transmitted. Normally the subscriber optical interface 140 with this traffic 1035 will transmit in a similar point during each superframe, but it is possible that another subscriber optical interface 140 will have requested and been granted a very long timeslot earlier in superframe₁ 1025. In this case, the subject traffic 1035 will not be transmitted until near the end of superframe₁ 1025. Inspection of Figure 10 will show that the total delay to traffic 1035 is slightly less than four superframes, represented by the queuing delay 1055. The total delay, from the time the traffic 1035 arrives at the subscriber optical interface, to the time when it has been completely transmitted to the transceiver node 120 is usually the total SOI delay 1065.

As described under Figure 9, the process illustrated in Figure 10 is a “sliding” process, in which the same event occurs every superframe. Thus, while the subject traffic 1035 is not being transmitted during superframes 1015 and 1020, other data is being handled during these times. Those skilled in the art will appreciate that the example of Figure 10 is a worst-case delay, and that average delays is lower.

Referring now to Figure 11, this figure illustrates an exemplary method for a handling upstream communications originating from subscribers of an optical network that are transmitted to the data service hub 110 of the optical network. Basically, Figure 11 provides an overview of the processing performed by transceiver node 120, and specifically, the optical tap routing device 435.

The description of the flow charts that follows is represented largely in terms of processes and symbolic representations of operations by conventional computer components, including a processing unit (a processor), memory storage devices, connected display devices, and input devices. Furthermore, these processes and

operations may utilize conventional computer components in a heterogeneous distributed computing environment, including remote file servers, computer servers, and memory storage devices. Each of these conventional distributed computing components can be accessible by the processor via a communication network.

5 The processes and operations performed below may include the manipulation of signals by a processor and the maintenance of these signals within data structures resident in one or more memory storage devices. For the purposes of this discussion, a process is generally conceived to be a sequence of computer-executed steps leading to a desired result. These steps usually require physical manipulations of physical
10 quantities. Usually, though not necessarily, these quantities take the form of electrical, magnetic, or optical signals capable of being stored, transferred, combined, compared, or otherwise manipulated. It is convention for those skilled in the art to refer to representations of these signals as bits, bytes, words, information, elements, symbols, characters, numbers, points, data, entries, objects, images, files, or the like.
15 It should be kept in mind, however, that these and similar terms are associated with appropriate physical quantities for computer operations, and that these terms are merely conventional labels applied to physical quantities that exist within and during operation of the computer.

It should also be understood that manipulations within the computer are often
20 referred to in terms such as creating, adding, calculating, comparing, moving, receiving, determining, identifying, populating, loading, executing, etc. that are often associated with manual operations performed by a human operator. The operations described herein can be machine operations performed in conjunction with various input provided by a human operator or user that interacts with the computer.

25 In addition, it should be understood that the programs, processes, methods, etc. described herein are not related or limited to any particular computer or apparatus. Rather, various types of general purpose machines may be used with the following process in accordance with the teachings described herein.

30 The logic flow described in Figure 11-16 can be the core logic or top level processing and can be executed repeatedly. The logic flow diagram illustrated in

Figures 11-16 show a process that can occur after initialization of the software or hardware components or both illustrated in Figures 1-5.

For example, in an object-oriented programming environment, software components or software objects or hardware that could be used to perform the steps illustrated in Figure 11-16 can be initialized or created. Therefore, one of ordinary skill in the art recognizes that several steps pertaining to initialization of software objects or hardware described in Figures 1 through 5 may not be illustrated.

The present invention may comprise a computer program or hardware or a combination thereof which embodies the functions described herein and illustrated in the appended flow charts. However, it should be apparent that there could be many different ways of implementing the invention in computer programming or hardware design, and the invention should not be construed as limited to any one set of computer program instructions. Further, a skilled programmer would be able to write such a computer program or identify the appropriate hardware circuits to implement the disclosed invention without difficulty based on the flow charts and associated description in the application text, for example. Therefore, disclosure of a particular set of program code instructions or detailed hardware devices is not considered necessary for an adequate understanding of how to make and use the invention. The inventive functionality of the claimed computer implemented process will be explained in more detail in the following description in conjunction with the remaining Figures illustrating the process flow.

Certain steps in the processes or process flow described below must naturally precede others for the present invention to function as described. However, the present invention is not limited to the order of the steps described if such order or sequence does not alter the functionality of the present invention. That is, it is recognized that some steps may be performed before or after other steps without departing from the scope and spirit of the present invention.

Briefly referring back to Figure 2, one transceiver node 120 can service a plurality of subscriber optical interfaces 140, by way of optical taps 130. Figure 3

illustrates more detail of the transceiver node 120, wherein a single optical tap routing device 435 works with a plurality of tap multiplexers 440, with each tap multiplexer 440 connected to a plurality of subscriber optical interfaces 140. Each tap multiplexer 440 can execute the process described in Figure 11. While two like processes are
5 illustrated in Figure 11, there may be six or more such processes operating simultaneously in a parallel manner. Only one of the plurality of processes will be described in detail below with reference to Figure 11.

Step 1205 is the first step in the method for processing upstream data packets as illustrated in Figure 11. In step 1205, during each superframe, each SOI 140 sends
10 a status message indicating how much data that SOI 140 has in its buffer to transmit. This amount of data is referred to as instantaneous queue size. After transmitting its status message, the SOI sends as much data as was authorized in that superframe (not shown on this flow chart). For example, during a superframe, typically each SOI transmits once and only once in one exemplary embodiment. During that
15 transmission time, the SOI typically first transmits its status message, followed by its traffic for the T_{min} (e.g., phone calls or DS1 channels), then transmits all other packets for which it was authorized.

In step 1205, the first SOI 140 sends its status message followed by its data. In step 1210, the tap multiplexer 440 determines if there are other SOIs 140 that have
20 not yet sent their status message. If all status messages have not been received during the current superframe, then the process loops back to step 1205 to wait for a status message from another SOI 140. In other words, the tap multiplexer 440 typically can't receive the next status message until the next SOI transmits, so it may wait. Note that this process occurs during superframe_{j,2} 910 as illustrated in Figure 9.

25 After all messages have been received, the process proceeds to step 1215, during which the transceiver node processes reservation information such as the instantaneous queue size reported by all SOIs in their status messages. The reservation information tells the tap multiplexer 440 how much data the particular SOI 140 has in its queue to transmit. This information will be used to determine the time

allocation for that SOI during superframe. This occurs during the RCD time 925 as illustrated in Figure 9.

After all the needed data is accumulated, in routine 1220, much of the upstream data management can be performed. Routine 1220 can comprise an emulated token bucket algorithm which can control the transmission of upstream data for multiple subscriber optical interfaces 140. While this routine 1220 has been characterized as a token bucket algorithm, routine 1220 can perform a little differently than a conventional token bucket as will be described in further detail below. Those skilled in the art are familiar with token bucket algorithms. One reference which describes such bucket algorithms is the following publication: "Policing and Shaping Overview," published by Cisco Systems, Inc., pages QC 87-QC 98. Another exemplary publication describing token bucket algorithms is the following white paper: "Cisco IOS(TM) Software Quality of Service Solutions," published by Cisco Systems, Inc., copyright 1998. The contents of both these reference are incorporated fully herein by reference.

At the end of Routine 1220, it has been determined when and for how long each SOI 140 will transmit during the next superframe. This information is transmitted downstream to the SOIs 140 during step 1225, which occurs during the RTD time 935 of Figure 9.

In step 1230, during superframe, 920 of Figure 9, the data that was assigned timeslots in routine 1220 is transmitted upstream, in step 1230. As data becomes available from the parallel processing of chains or groups of steps 1205 through 1230 performed by individual tap multiplexers 440, the information is passed upstream to data service hub 110, in step 1235.

Referring now to Figure 12, routine 1220 will now be explained in more detail. Processing for routine 1220 begins at step 1305. The subsequent steps are performed individually for each SOI 140 connected to a respective tap multiplexer 440. In step 1310, the queue size, that is, the amount of data that an individual SOIn 140n has to transmit, is evaluated. (The suffix n with reference to a subscriber optical interface 140 is used to indicate any one of a plurality of SOIs 140.)

In decision step 1310 it is determined whether the instant SOI_n 140n is active at the moment. If the inquiry to decision step 1310 is positive then it will be given its due allocation of time in which to transmit data upstream. However, if the inquiry to decision step 1310 is negative where the SOI 140 remains inactive for some consecutive number of superframes, then the bandwidth of the instant SOI 140 is taken away, and may be given to an active SOI 140 which can use that bandwidth. In the invention, some minimum bandwidth is usually reserved for each SOI 140 in order that certain housekeeping functions can be performed. These housekeeping functions can comprise monitoring the internal performance of the SOI 140 and learning quickly when an SOI 140 does have some data to transmit. The determination of whether or not the instant SOI_n 140 is active is made in the next few steps.

An activity counter for each SOI 140 is also usually provided, and is initially set to some maximum value determined by the network service provider or manufacturer of an SOI 140. If, by way of example, the maximum value is set equal to five then if the SOI 140 does not have any traffic to send for five consecutive superframes, it will be considered inactive.

In step 1310, it is also determined if the queue size of the instant SOI_n exceeds a minimum threshold that corresponds to its housekeeping traffic. If the inquiry to decision step 1310 is negative, then the "No" path is followed to step 1320 in which the activity counter is decremented. If the activity counter reaches zero, then bandwidth will be removed from the instant SOI being evaluated and can be assigned to other SOIs 140. If the inquiry to decision step 1310 is positive, where there is data in excess of the minimum queue size Q_{min} , then the "yes" branch is followed to step 1315 in which the activity counter is reset to its maximum value.

In routine 1325, the offered load (which is typically equal to the queue size less than what is to be transmitted during the current superframe_{j,1}) is filtered in the time domain. The purpose of this filter function in routine 1325 is to keep the system from over-reacting to instantaneous increases or decreases in load size. The problems

associated with over-reaction are understood by those skilled in the art. Further details of routine 1325 will be described below with respect to Figure 13.

Next, in routine 1330, the number of tokens in the emulated token bucket algorithm for $SOI_n 140$ is updated. Further details of routine 1330 will be described below with respect to Figure 14. This updated number of tokens will be used in routine 1335, which determines the maximum length of the time slot T_{maxn} that $SOI_n 140$ will be allowed to transmit data. Further details of routine 1335 will be described below with respect to Figure 15.

A minimum time slot, T_{minn} , comprising the minimum time that the $SOI_n 140$ will be allowed to transmit is determined in decision step 1340. If the inquiry to decision step 1340 is negative where the activity counter described above has reached zero, the "No" path is followed to step 1350 in which it is assumed that $SOI_n 140$ is inactive. The $SOI_n 140$ is then assigned the minimum time slot, which comprises a size just large enough to accommodate housekeeping functions.

If the activity counter of the $SOI_n 140$ is greater than zero, or if the $SOI_n 140$ has been granted some sustained rate that is constantly available regardless of whether it is used (such as if, a DS1 circuit was provided), then the "yes" branch is followed to step 1345. Those skilled in the art understand that DS1 circuits are often used by businesses to carry PBX (local telephone switching) and other data. Also, it may be that the user has some data that occurs sporadically, but which must be handled with minimum delay when the data is transmitted. A telephone call might fall in this category.

As noted previously if either of these two conditions exist, then the "Yes" path is followed to step 1345, in which T_{minn} is assigned an appropriate value for that $SOI_n 140$. The appropriate value is typically determined when service is established or modified according to subscription or level of service requested by the subscriber.

In decision step 1355, it is determined whether all $SOIs 140$ have been accounted for or processed during the instant superframe. If the inquiry to decision step 1355 is negative, the "No" branch is followed back to step 1310 to process data for the next $SOI 140$. If all $SOIs 140$ have been processed, then the "Yes" branch is

followed to step 1360, where the SOIs 140 are ordered or sorted by increasing Tmax. This sort determines the order that the SOIs will be considered in subsequent calculations, but it does not effect the order in which they transmit.

After the SOIs 140 are ordered in step 1360, the actual transmit times are
5 computed in routine 1365, which will be explained in more detail below in conjunction with Figure 16. After the transmit times are computed in routine 1365, the token bucket emulation ends and the process returns to step 1225.

Referring now to Figure 13, routine 1325 describing the updating of the offered load, will now be described. Figure 13 illustrates a filtering function that
10 reduces variation in the offered load, or queue size. This filtering is understood by those skilled in the art, to improve the performance of a communications system by preventing it from over-reacting to a sudden increase or decrease in traffic. The filter can use a large fraction of the previous load (which itself is the output of this routine) and a small fraction of the load that exists for the current superframe being evaluated,
15 to determine an average load.

In step 1410, the previous value of the average load is multiplied by a fraction f. Many values can be used for fraction f, and are within the scope of the invention, but a particularly beneficial and exemplary value is fifteen-sixteenths (15/16). Next, in step 1415, the instantaneous load of SOI_n 140 is computed. This load is the current
20 instantaneous queue size reported by SOI_n 140 during its last reporting time (during superframe_{j-2} 910), less the actual amount of data to be transmitted during the instant superframe_{j-1} 915. This value can be calculated because that transmission time allocation (to be used during the current superframe_{j-1} 915) has already been sent to SOI_n 140.

In step 1420 the result of this last computation is multiplied by the fraction (1-f). In step 1425 the result of the computations in steps 1410 and 1420 are added, and this value of step 1425 becomes the new offered load "L", to be used in future computations. In step 1430, this L value is stored, to be used in step 1410 the next time the subroutine is entered for this SOI_n. The subroutine ends at step 1435, and the
30 process returns to routine 1330 of Figure 12.

Routine 1330 comprises the updating of the token bucket, and is illustrated in detail in Figure 14. The update function begins in first step 1505. In step 1510, tokens are added to the token bucket for SOI_n 140 which is equivalent to the current SOI_n 's sustained rate times the duration of the superframe. By way of example, if the authorized sustained rate is 10 Mb/s and the duration of the superframe is 8 ms, the number of tokens added would correspond to 10 Mb/s times 8 ms, or 80,000 bits. This represents the number of bits that can nominally be transmitted during a superframe, assuming the SOI has them to transmit, and no other factors modify what it is allowed to transmit.

In step 1515, the number of bits corresponding to the amount of data transmitted in the previous superframe are subtracted, to account for what had been transmitted at that time. In decision step 1520, it is determined whether the token bucket now contains more tokens than are equivalent to the current SOI 's 140 burst limit. The inquiry to decision step 1520 may be positive if little had been transmitted in the previous superframe. If extra tokens were allowed to remain in the bucket, then the SOI 140 would be allowed to transmit more data than is allowed by the paid service level.

Thus, if there are more tokens than are equivalent to the burst limit, the inquiry to decision step 1520 yields a positive output, and the "yes" branch is followed to step 1525 where the extra tokens are removed, reducing the tokens in the bucket to the burst limit. If the inquiry to decision step 1520 is negative, then the "No" branch is followed to step 1530 where the subroutine ends.

Routine 1335 is another subprocess that can affect the amount of data that will be processed. This routine 1335 is shown in more detail in Figure 15. In routine 1335, it is determined whether the rate at which the SOI 140 is trying to transmit exceeds its allowed peak rate. In step 1610, the maximum timeslot to be granted to the SOI_m , T_{max} , is first set equal to SOI_n 's 140 sustained rate times the superframe duration (to compute how many bits the SOI 140 would be allowed to transmit in order to be transmitting at its allowed sustained rate), plus the current bucket allocation.

In decision step 1615, it is determined if this rate exceeds the allowed peak rate times the superframe duration. If the inquiry to decision step 1615 is positive, the "Yes" branch is followed to step 1620 in which the value of T_{maxn} is changed so that it equals the allowed peak rate (times the superframe duration). After step 1620, the algorithm ends at step 1625 where the process returns. If the inquiry to decision step 1615 is negative, when the value of T_{maxn} computed in step 1610 is less than the peak rate times the superframe duration, the "No" branch is followed to step 1625 where the algorithm ends. The process then returns to decision step 1340 of Figure 12.

Referring now to Figure 16, routine 1365 will be described below where the max-min algorithm that ultimately computes the amount of time each SOI 140 will be allowed to transmit in a subsequent superframe. Step 1705 is the first step of routine 1365. In step 1710, T_{maxn} for all SOIs 140 is added to obtain a total T_{max} . This can represent the amount of transmit time all SOIs 140 are requesting and to which they are entitled. The SOIs 140 will receive this transmit time if it does not exceed the transmit time available.

To determine how much transmit time is available, the total minimum transmit time, T_{min} , from step 1345 or step 1350 for all SOIs 140 is subtracted from the length of the superframe in step 1715. The difference is the time left over after all SOIs 140 have been allowed to transmit their data that is considered essential. (If in the subsequent superframe in which these computations will be used to discover a new SOI 140, then the time of that discovery timeslot usually must also be subtracted at step 1715.)

The following steps will typically apportion the transmit time available in the superframe after allowing for the minimum transmit time T_{min} of each SOI 140, and the discovery time if necessary. In step 1720, the ordering of SOIs 140 according to the maximum transmit time, T_{max} , calculated in step 1360 is used. The SOI(s) 140 with the least need for bandwidth are considered first, in step 1720. In decision step 1725, it is determined whether there is enough time in the superframe to give the transmit time evaluated in step 1720 to every SOI 140. If the inquiry to decision step

1725 is positive where there is enough time (a likely outcome on the first time through the loop that will be shown shortly), then the “Yes” branch is followed to step 1730, in which each of the SOIs 140 is granted the transmit time that the least needy SOI 140 is receiving.

- 5 Since the least needy SOI(s) 140 is satisfied in step 1730, it is removed from further consideration in step 1735. In step 1740, the amount of time available in the superframe is updated by subtracting the total time just granted. In decision step 1745, it is determined if there are any more SOIs 140 needing additional time. In early iterations, the answer to decision step 1745 is likely yes, so “Yes” branch is
- 10 followed back to step 1720.

- After a few iterations, the number of SOIs 140 being considered has been reduced by the SOI(s) 140 that have been satisfied. Again in step 1720, the SOI(s) 140 with the next least need for bandwidth are considered. In one exemplary iteration, the current SOI(s) 140 may have already been granted the transmit time of
- 15 the least needy SOI 140, but the remaining SOIs 140 may need more time for transmitting.

- Decision step 1725 checks to see if the remaining transmit time needed by this SOI 140 can be given to all remaining SOIs 140. If the inquiry to decision step 1725 is positive, the “Yes” branch is again followed to step 1730 in which all remaining
- 20 SOIs 140 are granted this additional transmit time or bandwidth. In step 1735, the SOI(s) 140 that was just satisfied is removed from further consideration, and in step 1740, the time just allocated is removed from the available time.

- The loop between steps 1720 and 1745 repeats until one of two events occur: Either all SOIs 140 eventually are granted the total time that they need when they
- 25 entered the subroutine, in which case decision step 1745 yields a “N” result and the subroutine ends. The alternative is that at some point there will not be enough time to allocate to all remaining SOIs 140, the time needed by the least needy of the remaining SOIs. In this case, the decision step 1725 yields a negative inquiry, and the “No” branch is followed to step 1750.

If the inquiry to decision step 1725 is negative, not all SOIs 140 are going to receive their requested and authorized bandwidth. It may be that the network service operator desires to allow one or more users to claim a disproportionate share of the remaining bandwidth. This can be accomplished by assigning a weighting factor to each SOI 140. The weighting factor is assigned when the service to a particular SOI 140 is established, but it is not used until this point in the process. When this point in the processing is reached, T_{max} , in step 1750 is multiplied by the pre-assigned weighting factor. This product is referred to as T_{maxwn} , where the subscript n can indicate that this is done for each SOI 140 involved (i.e., those SOIs 140 that have not been satisfied as to allowed transmit time).

In step 1755, the allocation to be given to each SOI 140 is computed. First, each SOI's T_{maxwn} is divided by the sum of all T_{maxw} for all units involved. This yields a fraction of the available time that the subject unit will have. The fraction is then multiplied by the time available, to yield the time the subject SOI can have. This is done in step 1755 for all SOIs 140 involved, so that at the end of step 1755, each SOI has its transmit time that is weighted according to its relative need and its pre-assigned weighting. The subroutine then terminates at step 1760, and the process returns to step 1370 of Figure 12 and then back to step 1225 of Figure 11, where the computed transmission times are communicated from the transceiver node 120 to each SOI 140.

An alternate embodiment not illustrated in Figure 16 would be to add some steps between steps 1715 and 1720. These additional steps could test among all the tap multiplexers 440 to determine if every tap multiplexer 440 would have enough bandwidth to transmit the data it was about to accept back up to the headend. To do this, at step 1715, time available from all tap multiplexers 440 is added and compared to the time available on the transmission path back to the headend via optical waveguide transceiver 430. If this number is lower than the needs of all the tap multiplexers 440, a step similar to step 1755 could be added between steps 1715 and 1720.

Upstream Transmission Scheduling

The transceiver node 120 manages the upstream link or data path to the data service hub 110 by allocating bandwidth to each subscriber optical interface 140. This section of the detailed description describes that allocation process, and it illustrates several example configurations. Similar to downstream traffic, each channel that comprises an upstream data path of predetermined number of subscribers operates independently.

The transceiver node 120 manages upstream transmission by emulating a token bucket algorithm. Within the transceiver node 120, each subscriber optical interface 140 is typically provisioned with the following parameters:

- Peak Rate. The peak rate is the maximum rate that a subscriber is allowed to transmit upstream packets.

- Sustained Rate. The sustained rate is the profile rate for the subscriber. Normally this rate represents both the minimum bandwidth that the subscriber is assured of receiving and the maximum bandwidth that the subscriber can use over a long time period.

- Burst Size. The amount of data that the subscriber can transmit at its peak rate during bursts. A special value indicates that the subscriber has no limits on its burst size.

- Priority Weighting. The relative importance of this subscriber's data compared to other subscribers. The transceiver node 120 uses this value to arbitrate upstream transmissions when multiple subscriber optical interfaces 140 are attempting to transmit traffic bursts, the sum of which would exceed the link or upstream data path capacity.

- Guaranteed Indication. An indication of whether or not the subscriber must always be guaranteed its sustained rate. If true, then the transceiver node 120 usually must allocate the indicated bandwidth for the subscriber, even when that subscriber is not currently active. This setting lets a subscriber begin transmitting at its sustained rate immediately after becoming active. If this indication is false, then, when the

subscriber is idle, the transceiver node 120 may “borrow” bandwidth and allocate it to other subscriber optical interfaces 140.

The method of Figures 11-17 can have several key properties.

•In one exemplary embodiment, the minimum allocation to each subscriber optical interface 140 can be either its sustained rate or, if that subscriber’s bandwidth is not guaranteed, 256 Kbit/s.

•In one exemplary embodiment, the maximum allocation to each subscriber optical interface 140 can be either its peak rate or its current token bucket’s worth of burst, whichever is smaller.

•In one exemplary embodiment, if all subscriber optical interfaces 140 cannot be granted their maximum, then bandwidth over and above the minimum is allocated using a weighted max-min algorithm. The weight factor is the product of a provisioned weight and a low-pass filtered measurement of buffer occupancy of each subscriber optical interface 140. The max-min constraint for each subscriber optical interface 140 can be its maximum allocation.

As an example of the method illustrated in Figures 11-17, the Table 2 presents a simplified example with three subscribers. The first column shows each subscriber’s calculated limit; the second column shows the average offered load, and the third column indicates each subscriber’s priority weighting.

Table 2 - Subscriber Group comprising three subscribers

	Limit	Average Load	Weight	Allocation
SOI 1	500 Mbit/s	100 Kbit	1	128 Mbit/s
SOI 2	500 Mbit/s	150 Kbit	2	362 Mbit/s
SOI 3	10 Mbit/s	5 Kbit	1	10 Mbit/s

according to the last column in the table. Since the third subscriber optical interface 140 in Table 2 has the lowest limit, it receives a full grant of 10 Mbit/s. (The other two subscriber optical interfaces 140 also receive 10 Mbit/s.) The third subscriber optical interface 140, in this exemplary embodiment, is the only subscriber that can be allocated its full limit, however, so the remaining bandwidth (470 Mbit/s) is divided among the remaining two subscribers. The amount each subscriber receives can be proportional to the product of its load and its weight. The second subscriber optical interface 140, consequently, receives 75% of the 470 Mbit/s in addition to the 10 Mbit/s. The first subscriber optical interface 140 receives 25% of 470 Mbit/s plus 10 Mbit/s.

Subscriber Optical Interface 140 Packet Processing

The subscriber optical interface 140 supports Quality of Service Management of upstream traffic through a two-priority queuing mechanism. The subscriber optical interface 140 can comprise a standard layer-2 Ethernet switch; the last mile link or upstream data path can be treated simply as another Ethernet port for the switch. To provide priority for voice and TDM over IP traffic, the subscriber optical interface 140 switch is configured to give priority to packets generated by internal applications. After prioritization, upstream packets are queued in the output buffer until they can be transmitted to the transceiver node 120. The output buffer provides first-in first-out service. The size of the buffer is provisioned separately for each subscriber optical interface 140.

BACKBONE NETWORK INTEGRATION

Quality of service (QoS) is most powerful when it can be managed globally across an entire network, and the present invention provides unparalleled opportunities for global QoS management across an entire backbone network. The basis for this integration is IP's differentiated services (diffserv) architecture.

Creating Service Level Agreements

The transceiver node (TN) 120 provides extensive support for managing service level agreements (SLAs) with subscribers. Although the TN 120 is necessarily only one component in an overall agreement, as the access network, it is critical. The following examines how the last mile link or upstream data path contributes to SLAs and how the QoS management functionality can support SLAs.

Components of an SLA

Service level agreements are typically more common with private network technologies such as ATM or Frame Relay. The power and flexibility of the TN's 120 QoS management, however, permits those same concepts to be extended to IP access networks. The same components that are part of traditional ATM or Frame Relay SLAs can be part of an data path SLA.

- Peak Rate. The maximum rate at which the network will accept traffic bursts from the user, expressed in bits per second. The network discards traffic that exceeds the peak rate.
- Sustained Rate. The minimum throughput that the network will provide to the user, expressed in bits per second.
- Burst Size. The amount of traffic that the network will accept without pause at the user's peak rate, expressed in bits.
- Maximum Latency. The worst-case latency the user's traffic will experience as it traverses the network.
- Loss Rate. The percentage of traffic conforming to the peak rate, sustained rate, and burst size that the network may discard.

Of course, service providers can include other elements in their service level agreements. The transceiver node 120 provides a wealth of features that a service provider may position as value-added services. The TN 120 supports services such as the following:

- Application Prioritization. Giving priority to key network applications (e.g. Virtual Private Network traffic).
- Enhanced Statistics. Providing detailed traffic profiles and statistics to assist the user in network growth planning.
- 5 • Active Monitoring. Continuously monitoring user traffic to provide early detection of network application faults (e.g. Web server failures).
- Network Security. Providing encryption of traffic to the subscriber.

This part of the detailed description focuses on traditional SLA performance metrics. It examines how the Transceiver node 120 contributes to network performance, and how to provision downstream QoS management to meet SLA requirements. Table 3 below lists key parameters and values used in equations throughout this part of the description.

Table 3 - Inherent Link Characteristics

C	Link Capacity (500 Mbit/s)
τ	Superframe Period (8 ms)

15 Table 4- Upstream Configuration parameters (per Subscriber)

Bu	Upstream Burst Size (bit)
Pu	Upstream Peak Rate (bit/s)
Ou	SOI Output Buffer Size (bit)
Ru	Upstream Sustained Rate (bit/s)
Wu	Upstream Weight (unit-less)

Rigorous SLAs and Oversubscription

Because business requirements differ among service providers and among subscribers, the transceiver node 120 allows typically providers significant flexibility in enforcing SLA performance metrics. Some deployments can require ironclad service level agreements; those environments require a conservative provisioning strategy. Conservative provisioning can provide extremely tight performance guarantees, but it generally results in a lower overall network utilization and, ultimately, greater capital expenditures.

In other deployments (residential Internet access, for example) SLAs are not common and may not be desirable. In those environments a more aggressive provisioning strategy may be effective. In general, meaningful SLAs are usually not enforceable when a network is provisioned aggressively; the resulting networks, however, may be operated at much higher utilization.

This part of the description considers both strict SLAs and slightly relaxed SLAs. Relaxed SLAs allow a modest amount of oversubscription of network resources; in exchange, the service provider cannot offer rigorous guarantees for all aspects of network performance.

Upstream Performance

For strict SLA environments, service providers can offer tight guarantees on upstream performance, including peak rate, sustained rate, burst size, latency, and loss. To do that, the provider must adhere to the following constraints.

- The sum of the peak rates for all subscribers must be less than the link capacity. $[\sum P_U < C]$

- The burst size for each subscriber must be equal to the output buffer size configured in that subscriber's Last Mile Gateway. $[B_U = O_U \forall \text{ SOIs}]$

Under these conditions, it can be straightforward to relate transceiver node configuration parameters to SLA metrics.

Table 4 - Transceiver Node 120 Configuration Parameters to SLA Metrics

SLA Metric	Node Configuration Parameters
Peak Transmission Rate	equal to Upstream Peak Rate $[= P_U]$
Sustained Transmission Rate	equal to Upstream Sustained Rate $[= R_U]$
Transmission Burst Size	equal to Upstream Burst Size and SOI Output Buffer Size $[= B_U]$
Node Upstream Latency	equal to Upstream Burst Size divided by Upstream Sustained Rate plus superframe period $[= B_U/R_U + \tau]$
Node Upstream Loss Rate	0

A more lenient deployment could relax the peak rate constraint and replace it with a constraint on sustained rates. In this environment, the constraints are as follows.

- The sum of the sustained rates for all subscribers can be less than the link or upstream data path capacity.
[$\Sigma R_U < C$]

- The burst size for each subscriber can be set equal to the output buffer size configured in that subscriber's optical interface 140. [$B_U = O_U \forall \text{ SOIs}$]

- In this configuration, the worst-case peak transmission rate for a user may be less than its configured upstream peak rate. The worst-case upstream latency and loss rates also increase. Exact closed-form equations for SLA metrics are not possible in this configuration. Table 5 provides bounds on the SLA metrics that are suitable for SLA contracts.

Table 5 - Bounds on SLA Metrics Suitable for SLA Contracts

SLA Metric	Node Configuration Parameters
Peak Transmission Rate	either the Upstream Peak Rate or at least the weighted share of excess link capacity (capacity above the Upstream Sustained Rates of all LMGs, whichever is smaller [$\geq \min(P_U, R_U + (C - \Sigma R_U) \cdot B_U \cdot W_U / \Sigma(B_U \cdot W_U))$]
Sustained Transmission Rate	equal to Upstream Sustained Rate [$= R_U$]
Transmission Burst Size	equal to Upstream Burst Size and SOI Output Buffer Size [$= B_U$]
Node Upstream Latency	no more than the value calculated by the following iterative equations: $I_0 \leftarrow (P_U - R_U) \cdot \tau$; $A_0 \leftarrow 0$; $S_0 \leftarrow R_U \cdot \tau$ $I_{i+1} \leftarrow I_i - S_i$; $A_{i+1} \leftarrow 15/16 \cdot A_i + 1/16 \cdot I_i$; $S_{i+1} \leftarrow R_U \cdot \tau + A_i \cdot W_U \cdot (C - \Sigma R_U) / \Sigma(B_U \cdot W_U)$ $\text{delay} \leq i \cdot \tau \mid \Sigma S_i \geq (P_U \cdot \tau)$
Node Upstream Loss Rate	no more than the amount by which the latency exceeds the rigorous case [$\leq (\text{Latency} - B_U/R_U - \tau) / \text{Latency}$]

- It should be understood that the foregoing relates only to illustrate the embodiments of the present invention, and that numerous changes may be made therein without departing from the scope and spirit of the invention as defined by the following claims.